

Predicting Inflation: Professional Experts Versus No-Change Forecasts

Tilman Gneiting, Thordis L. Thorarinsdottir
Institut für Angewandte Mathematik
Universität Heidelberg, Germany

October 13, 2010

Abstract

We compare forecasts of United States inflation from the Survey of Professional Forecasters (SPF) to predictions made by simple statistical techniques. In nowcasting, economic expertise is persuasive. When projecting beyond the current quarter, novel yet simplistic probabilistic no-change forecasts are equally competitive. We further interpret surveys as ensembles of forecasts, and show that they can be used similarly to the ways in which ensemble prediction systems have transformed weather forecasting. Then we borrow another idea from weather forecasting, in that we apply statistical techniques to postprocess the SPF forecast, based on experience from the recent past. The foregoing conclusions remain unchanged after survey postprocessing.

Key words and phrases: inflation; predictive distribution; reference forecast; statistical postprocessing; Survey of Professional Forecasters.

1 Introduction

A wealth of societal decisions can benefit from accurate forecasts of future inflation, ranging from the setting of monetary and fiscal policies to negotiations of wage contracts and investment judgments. To predict inflation rates, various methods have been employed, including statistical time series techniques, methods based on term structures and yield

curves, and survey-based measures from consumers or professional experts. A prominent recent study argues that survey forecasts perform best (Ang et al., 2007).

The Survey of Professional Forecasters (SPF) is the leading quarterly survey of macroeconomic variables in the United States (Zarnowitz, 1969; Croushore, 1993). It began in 1968 and was conducted by the American Statistical Association and the National Bureau of Economic Research, before the Federal Reserve Bank of Philadelphia took over in 1990. The panel comprises university professors as well as private sector economists, who are asked each quarter to predict a range of macroeconomic variables for the current and each of the following four quarters. Our interest here is in inflation, in the form of the quarter-over-quarter change of the consumer price index expressed in annualized percentage points, which we now refer to as an inflation rate. Figure 1 displays observed United States inflation rates since the third quarter of 1995 along with the SPF forecasts issued a year earlier. Following common practice, we talk of a prediction horizon of one quarter when referring to current quarter nowcasts, and to prediction horizons of two to five quarters for the following four quarters (Croushore and Stark, 2001). For example, forecasts issued a year earlier correspond to a prediction horizon of five quarters.

The SPF panel forecast can be summarized to provide a single point forecast, where we also follow common practice and take it to be the median of the individual experts' predictions (Stark, 2010). The traditional no-change forecast equals the most recent available observed rate. This is the classical reference forecast in the economic literature and that used by the Federal Reserve Bank of Philadelphia. Here, we introduce a novel kind of simplistic reference forecast, which we call the probabilistic no-change forecast. It considers a rolling training period, consisting of the 20 most recent observed inflation rates, and takes the median thereof as a point predictor. This can be interpreted as a theoretically optimal point forecast under a white noise assumption for the inflation rates, whereas the traditional no-change forecast can be interpreted as optimal under a random walk model.




Figure 1: SPF forecasts of United States inflation at a prediction horizon of five quarters. The small blue dots show the individual SPF experts' point forecasts, with the pale blue boxes indicating their range. The large purple dots represent the observed inflation rates. Note the extreme deflation in the fourth quarter of 2008.

In Section 2 we compare the predictive performance of the SPF point forecast to the simple no-change forecasts. Essentially, in current quarter nowcasts economic expertise is persuasive. At prediction horizons beyond the current quarter, the SPF forecast outperforms the traditional no-change forecast, but not the novel, equally simplistic probabilistic no-change forecast.

The SPF forecast can also be interpreted as an ensemble forecast, similar to the ways in which weather and climate scientists have been using ensemble prediction systems with great success (Palmer, 2002). State of the art weather forecasting uses ensembles whose members are point forecasts from numerical weather prediction models, with the members differing in initial conditions and/or the specifics of the numerical model used. For ensemble weather forecasts, some form of statistical postprocessing is required to correct for model biases and insufficient representations of the forecast uncertainty (Gneiting and Raftery, 2005). From Figure 1, we see that similar to weather ensembles the SPF ensemble forecast is uncalibrated, in that too many observations fall outside the range of the ensemble forecast.

Postprocessing methods in meteorology provide statistically corrected predictive probability distributions for future weather quantities that condition on the ensemble forecast. From the predictive distribution, the probability of any event of interest can be computed, and one can issue the optimal point forecast under the loss function at hand (Diebold et al., 1998; Engelberg et al., 2009). We adopt this approach and develop statistical postprocessing methods for the SPF by using heteroscedastic regression and Gaussian mixture models. In Section 3 these methods are introduced in detail and their predictive performance is evaluated, with results that resemble those without postprocessing. Finally, in Section 4 we study the robustness of our results. Concluding remarks are given in Section 5, where we discuss the methodological as well as the economic and societal implications of our work.

2 Predictive performance: SPF versus no-change forecasts

The Survey of Professional Forecasters (SPF) panel comprises academic as well as private sector economists, who are asked to provide point forecasts for a range of macroeconomic variables. They are also requested to provide probability forecasts, but these refer to annual, rather than quarterly, percentage change, and thus are not considered here. The forecasts are issued in the middle of a quarter for the current quarter (a prediction horizon of one quarter) and each of the following four quarters (prediction horizons of two quarters through five quarters). While Engelberg et al. (2009) show that the predictions are mostly consistent with the hypothesis that SPF panel members report their subjective means, medians or modes, they also note that SPF forecasters tend to give more favorable views of the economy than warranted by their subjective probabilities.

Inflation rates based on the consumer price index (CPI) have been included in the survey since the third quarter of 1981, and forecast data are available on-line at <http://www.phil.frb.org/econ/spf/>. As realizing observations, we use the CPI vintage from the Real-Time Data Research Center of the Federal Reserve Bank of Philadelphia

(Croushore and Stark, 2001), which is available from the third quarter of 1994 on. The traditional and probabilistic no-change forecasts then use the most recent vintage available at the issuing time of the forecast, while the predictive performance is evaluated against the May 2010 vintage. To convert the original monthly observations of the CPI into annualized quarterly growth rates, we follow common practice by averaging the monthly observations of each quarter, and using the formula

$$y_t = \left(\left(\frac{z_t}{z_{t-1}} \right)^4 - 1 \right) \times 100,$$

where z_t is the observed quarterly CPI in quarter t , and y_t is the observed quarter-over-quarter growth rate of the CPI, or simply, the inflation rate in quarter t , in percentage points.

We now compare the predictive performance of the SPF median forecast to the traditional and probabilistic no-change forecasts. Due to the dynamic data vintage, the traditional no-change forecast does not exactly trail the observations, even though it does so approximately. The probabilistic no-change forecast uses the median of the 20 most recent observations available at the issuing time. Table 1 summarizes the predictive performance for the period from the third quarter of 1995 through the first quarter of 2010 in terms of the mean absolute error (MAE). In addition, we report the results of Diebold and Mariano (1995) tests of the hypothesis of equal predictive performance between the SPF forecast and the reference forecasts. In doing so, we provide the lower tail probability under the null hypothesis in percentage points, where a value of 00 indicates a lower tail probability less than or equal to 1%, a value of 01 a lower tail probability between 1% and 2%, ..., and a value of 99 a lower tail probability exceeding 99%. Thus, values from 00 to 04 and 95 to 99 correspond to a statistically significant difference at the 5% level for a one-sided test, and at the 10% level for a two-sided test.

The table confirms the well known fact that the SPF median forecast outperforms the

Table 1: Mean absolute error (MAE) for point forecasts of United States inflation from the third quarter of 1995 to the first quarter of 2010, in percentage points, along with the lower tail probability of the Diebold-Mariano test for equal predictive performance between the SPF forecast and the reference forecast.

Forecast method	Prediction horizon in quarters				
	1	2	3	4	5
SPF	0.89	1.44	1.51	1.49	1.49
Probabilistic no-change	1.45 ₉₉	1.46 ₅₈	1.45 ₁₈	1.48 ₃₉	1.48 ₄₄
Traditional no-change	1.81 ₉₉	2.06 ₉₆	2.00 ₉₁	2.06 ₉₅	2.03 ₉₇

traditional no-change forecast at all prediction horizons. The corresponding MAE differentials are statistically significant as well as substantial in size, reaching values close to a percentage point. At a nowcasting prediction horizon of a single quarter, the SPF median forecast also has substantially lower MAE than the probabilistic no-change forecast, even though the latter is more competitive than the traditional no-change forecast. However, at prediction horizons from two to five quarters, the SPF forecast is unable to outperform the probabilistic no-change forecast, showing MAE values that are about equal for the two methods.

Thus far, we have considered point forecasts, as opposed to probabilistic forecasts or predictive distributions, which are of ever increasing importance in a wide trans-disciplinary range of applications (Timmermann, 2000; Gneiting, 2008). The SPF and the probabilistic no-change method provide predictive distributions in natural ways, in that they can be identified with the discrete probability measures that assign equal mass to each of the experts, or each of the CPI observations in the training period, respectively. For example, the predictive distribution that corresponds to our standard version of the probabilistic no-change forecast assigns mass $1/20$ to each of the 20 most recent inflation observations available at the issuing time. To obtain a predictive distribution associated with the traditional no-change forecast, we take it to be Gaussian, with mean equal to the most recent

Table 2: Mean continuous ranked probability score (CRPS) for probabilistic forecasts of United States inflation from the third quarter of 1995 to the first quarter of 2010, in percentage points, along with the lower tail probability of the Diebold-Mariano test for equal predictive performance between the SPF forecast and the reference forecast.

Forecast method	Prediction horizon in quarters				
	1	2	3	4	5
SPF	0.69	1.16	1.25	1.26	1.27
Probabilistic no-change	1.08 ₉₉	1.10 ₂₁	1.10 ₀₀	1.10 ₀₀	1.11 ₀₀
Traditional no-change	1.56 ₉₉	1.66 ₉₇	1.51 ₈₂	1.57 ₉₀	1.47 ₈₇

available observation, and variance equal to the empirical mean squared error (MSE) of the traditional no-change forecast over the rolling 20-quarter training period.

To assess the predictive performance of the probabilistic forecasts, we use the continuous ranked probability score (CRPS), which is a decision theoretically coherent proper scoring rule, and reduces to the absolute error in the case of a point forecast (Matheson and Winkler, 1976; Gneiting and Raftery, 2007). If the predictive cumulative distribution function (CDF) is F and the observation y verifies, the CRPS is defined as

$$\text{CRPS}(F, y) = \int_{-\infty}^{\infty} (F(x) - \mathbb{1}\{x \geq y\})^2 dx, \quad (1)$$

where $\mathbb{1}\{x \geq y\}$ denotes an indicator function that attains the value 1 if $x \geq y$ and the value 0 otherwise. Grit et al. (2006) showed that for a discrete probability measure F that puts mass $1/M$ on each of x_1, \dots, x_M , this can be written as

$$\text{CRPS}(F, y) = \frac{1}{M} \sum_{m=1}^M |x_m - y| - \frac{1}{M^2} \sum_{m=1}^M \sum_{n=1}^M |x_m - x_n|.$$

Table 2 summarizes the predictive performance of the SPF, probabilistic no-change and traditional no-change forecasts in terms of the mean CRPS over the foregoing test period. Again, the SPF forecast outperforms the extant reference standard, namely the traditional




Figure 2: Mean absolute error (MAE, red solid line) and mean continuous ranked probability score (CRPS, red dashed line) for five-quarter ahead probabilistic no-change forecasts for United States inflation from the third quarter of 1995 to the first quarter of 2010, as functions of the length of the rolling training period. For comparison, we also show the MAE (blue solid line) and the mean CRPS (blue dashed line) for the SPF forecast.

no-change forecast, at all prediction horizons. In nowcasts for the current quarter, the SPF forecast also shows substantially lower CRPS than the probabilistic no-change forecast. Beyond the current quarter in true forecasting mode the simplistic probabilistic no-change forecast outperforms the SPF experts. Furthermore, as Figure 2 demonstrates, these results are robust to the choice of the length of the rolling training period for the probabilistic no-change forecast.

A valid concern at this point is that the SPF panel is simply a collection of point forecasts that is not necessarily meant to be taken as a discrete predictive distribution. As noted, this resembles the situation in weather forecasting where meteorologists use ensembles of point forecasts from numerical weather prediction models (Palmer, 2002), which are subject to biases and dispersion errors, thus calling for statistical postprocessing (Gneiting and Raftery, 2005). Indeed, the Federal Reserve Bank of Philadelphia has recently experimented with Gaussian density forecasts that derive from the SPF panel (Stark, 2010) and can be interpreted as postprocessing methods. In the next section, we take up this idea and develop statistical postprocessing techniques that are tailored to the SPF panel.

3 Survey postprocessing

In weather forecasting, statistical postprocessing methods have been used with great success to improve the predictive performance of ensemble prediction systems, with heteroscedastic regression (Gneiting et al., 2005; Thorarinsdottir and Gneiting, 2010) and Bayesian model averaging (Raftery et al., 2005; Sloughter et al., 2010) being state of the art techniques. Here we introduce variants of these methods that are tailored to the SPF ensemble and inflation forecasts. In the SPF, the composition of the expert panel changes gradually over time, with individual members providing forecasts for about six years on average (Engelberg et al., 2009). Furthermore, there are missing forecasts for essentially all members even during the period when they are active. Thus, it is very difficult to assess the predictive performance of individual members, in contrast to what is commonly done in weather forecasting. Therefore, the predictive distributions obtained here depend on summary statistics derived from the SPF panel or the probabilistic no-change forecast. In this context, we denote the median and the variance of the SPF panel by μ_{SPF} and σ_{SPF}^2 . Similarly, we write μ_{PNC} and σ_{PNC}^2 for the median and the variance of the probabilistic no-change (PNC) forecast with a rolling 20 quarter training period. The density function and the CDF of the standard normal distribution will be denoted by φ and Φ , respectively.

3.1 Heteroscedastic regression

Gneiting et al. (2005) and Thorarinsdottir and Gneiting (2010) proposed a statistical postprocessing method for ensembles of point forecasts that uses heteroscedastic regression (HR) (Leslie et al., 2007), where the location parameter of the predictive distribution is a linear function of the ensemble member forecasts, and the scale parameter is a linear function of the ensemble variance. Here, we adapt the method so that the postprocessed predictive distribution is the asymmetric three-parameter two-piece normal distribution (John,

1982) with density

$$f_{\text{TPN}}(y) = \begin{cases} \left(\frac{2}{\pi}\right)^{1/2} (\sigma_1 + \sigma_2)^{-1} \exp\left(-\frac{(y - \mu)^2}{2\sigma_1^2}\right) & \text{if } y \leq \mu, \\ \left(\frac{2}{\pi}\right)^{1/2} (\sigma_1 + \sigma_2)^{-1} \exp\left(-\frac{(y - \mu)^2}{2\sigma_2^2}\right) & \text{if } y \geq \mu, \end{cases}$$

and CDF

$$F_{\text{TPN}}(y) = \begin{cases} \frac{2\sigma_1}{\sigma_1 + \sigma_2} \Phi\left(\frac{y - \mu}{\sigma_1}\right) & \text{if } y \leq \mu, \\ \frac{\sigma_1 - \sigma_2}{\sigma_1 + \sigma_2} + \frac{2\sigma_2}{\sigma_1 + \sigma_2} \Phi\left(\frac{y - \mu}{\sigma_2}\right) & \text{if } y \geq \mu. \end{cases} \quad (2)$$

The two-piece normal distribution has been used for density forecasts of United Kingdom inflation by the Bank of England since 1996 (Wallis, 1999; Elder et al., Autumn 2005). Its asymmetry allows for a distinct treatment of upside and downside risks in inflation forecasting. If $\sigma_1 < \sigma_2$ the distribution is right skewed, and both the mean and the median exceed its mode, μ . If $\sigma_1 > \sigma_2$ the distribution is left skewed.

We consider two variations of the HR approach. The first variant models the parameters of the two-piece normal distribution as functions of the SPF median and SPF variance, in that

$$\mu = a + b\mu_{\text{SPF}}, \quad \sigma_1^2 = c_1 + d_1\sigma_{\text{SPF}}^2, \quad \sigma_2^2 = c_2 + d_2\sigma_{\text{SPF}}^2, \quad (3)$$

and we refer to it as the HR model with SPF covariates. The second nests the first and posits

$$\mu = a + b_1\mu_{\text{SPF}} + b_2\mu_{\text{PNC}}, \quad \sigma_1^2 = c_1 + d_{11}\sigma_{\text{SPF}}^2 + d_{12}\sigma_{\text{PNC}}^2, \quad \sigma_2^2 = c_2 + d_{21}\sigma_{\text{SPF}}^2 + d_{22}\sigma_{\text{PNC}}^2. \quad (4)$$

We refer to the specification in (4) as the HR model with SPF and PNC covariates. In out-of-sample forecasting overfitting is heavily penalized, and it is an empirical question whether or not this more complex model can outperform the more parsimonious one.

Table 3: Minimum CRPS estimates for the HR density forecast (4) of United States inflation in the first quarter of 2008 at a prediction horizon of two quarters, using a rolling training period of 40 quarters. The resulting parameter values for the two-piece normal distribution (2) are mean value of $\mu = 1.90$, with standard deviations $\sigma_1 = 0.59$ and $\sigma_2 = 3.27$, respectively.

a	b_1	b_2	c_1	d_{11}	d_{12}	c_2	d_{21}	d_{22}
0.36	0.53	0.00	0.00	0.52	0.00	0.00	0.00	3.10

In Gneiting et al. (2005) and Thorarinsdottir and Gneiting (2010), the parameters of the postprocessing model were estimated by minimizing the mean CRPS over the training period, and this was shown to yield slightly better predictive performance than maximum likelihood estimation. In an economic context, similar approaches have been discussed by Elliott and Timmermann (2008). We adopt this proposal, using formula (8) in the Appendix, and performing the minimization numerically, via the Broyden-Fletcher-Goldfarb-Shanno algorithm as implemented in R (R Development Core Team, 2008). Each prediction horizon requires its individual fit with distinct parameter estimates, using a rolling training period from the recent past.

As an example, Figure 3 and Table 3 illustrate the HR density forecast with SPF and PNC covariates for the first quarter of 2008 at a prediction horizon of two quarters, fitted on a rolling training period of 40 quarters. The point forecasts of the SPF experts had median 2.90 and standard deviation 0.82. The PNC ensemble had a higher median, at 3.30, and a higher standard deviation, at 1.86. Table 3 shows the parameter estimates for the HR model (4). The mode μ of the two-piece normal distribution (2) is determined by the SPF median, the downside risk σ_1^2 by the SPF spread, and the upside risk σ_2^2 by the PNC spread. Figure 3 shows the SPF and PNC ensembles along with the postprocessed density forecast, which is strongly right skewed, with mode at 1.90 and median at 3.67, higher than both the SPF median and the PNC median. The verifying inflation rate in the first quarter of 2008 was 4.66 percentage points.




Figure 3: Density forecast for the United States inflation rate in the first quarter of 2008 at a prediction horizon of two quarters, using the HR model (4) with both SPF and PNC covariates. At bottom, the SPF panel (blue bars), which has median 2.90, and the PNC ensemble (red bars), with median 3.25, are shown. The median of the density forecast is indicated by the solid orange line at 3.67, and the realizing inflation rate is represented by the purple dashed line at 4.66.

3.2 Gaussian mixture models

Bayesian model averaging (BMA) is a standard method for combining inferences (Hoeting et al., 1999). Its use in the statistical postprocessing of ensemble weather forecasts was proposed by Raftery et al. (2005) and Sloughter et al. (2010). In the normal mixture version of Raftery et al. (2005), the BMA predictive density is a mixture of Gaussian densities, where the components are associated with individual ensemble member, and the mixture weights reflect the members' relative contribution to predictive skill over the training period.

Here we use a similar idea, taking the postprocessed predictive distribution to be a mixture of two Gaussian components, with a CDF of the form

$$F_{\text{GM}}(y) = \alpha \Phi\left(\frac{y - \mu_1}{\sigma_1}\right) + (1 - \alpha) \Phi\left(\frac{y - \mu_2}{\sigma_2}\right). \quad (5)$$

In a first variant we put

$$\mu_1 = \mu_{\text{SPF}}, \quad \sigma_1 = \sigma_{\text{SPF}}, \quad \mu_2 = \mu_{\text{PNC}}, \quad \sigma_2 = \sigma_{\text{PNC}}, \quad (6)$$

so that the mixture weight $\alpha \in [0, 1]$ is the only parameter to be estimated. This is our most parsimonious, standard Gaussian mixture (GM) model. In a second variant, we put

$$\mu_1 = \mu_{\text{SPF}}, \quad \mu_2 = \mu_{\text{PNC}}, \quad (7)$$

with $\alpha \in [0, 1]$, $\sigma_1 > 0$ and $\sigma_2 > 0$ to be estimated. We refer to (7) as the GM model with variance adjustment. In our experience, bias correction of the location parameters deteriorates the predictive performance out of sample, and thus we do not present results for such models. The parameters are estimated by maximum likelihood via the expectation-maximization (EM) algorithm, based on training data from a rolling training period, as described by Raftery et al. (2005) and implemented in the R package `ENSEMBLEBMA`.

An example of a GM predictive distribution with variance adjustment is given in Figure 4, using a rolling training period of 40 quarters.

3.3 Predictive performance

In the SPF record, quarterly forecasts and observations of the CPI are available for 115 quarters. We split the data such that the forecasts through the second quarter of 1995 are used solely for training purposes, while the methods are tested on the data thereafter. The CPI vintage data have been released quarterly since the third quarter of 1994 and it would thus not have been possible to issue real-time forecasts based on this data prior to the third quarter of 1994. For parameter estimation, we use a rolling training period of 40 quarters. The PNC ensembles used are based on 20 quarterly observations.

Figure 5 illustrates the postprocessed predictive distributions using the HR method with




Figure 4: Density forecast for the United States inflation rate in the second quarter of 2005 at a prediction horizon of two quarters (in pale green), using the GM model (7) with variance adjustment. At bottom, the SPF panel (blue lines), which has median 2.20 and standard deviation 0.85, and the PNC ensemble (red lines), with median 3.05 and standard deviation 1.39, are shown. The median of the density forecast is indicated by the green solid line at 2.49, and the realizing inflation rate is represented by the purple dashed line at 2.73. The corresponding parameter estimates are $\alpha = 0.59$, $\sigma_1 = 0.98$ and $\sigma_2 = 1.30$. The densities corresponding to the SPF panel and the PNC forecast are indicated accordingly.

SPF and PNC covariates, and the GM technique with variance adjustment, for the period from 2006 on. Through 2009, the methods yield very similar predictive distributions, while their spread diverge for the period from early 2010 on. This goes hand-in-hand with the growing discrepancy between the median of the SPF and the PNC ensembles.

In the performance comparison below, we include a probabilistic forecast with conditional predictive CDF given by

$$F(y) = \Phi\left(\frac{y - \mu_{\text{SPF}}}{\text{MSE}_{\text{SPF}}}\right)$$

where MSE_{SPF} denotes the mean squared error of the SPF median over the past 40 quarters. The Federal Reserve Bank of Philadelphia has been experimenting with this method (Stark,




Figure 5: Recent forecasts of the United States inflation rate at a prediction horizon of five quarters. The point forecasts shown are the medians from the SPF panel data (blue line), the PNC ensemble (red line), the HR method with SPF and PNC covariates (orange line), and the GM technique with variance adjustment (green line). The 80% prediction intervals for the HR method are indicated in orange, the 80% prediction intervals for the GM technique are indicated in pale green, and the large purple dots represent the observed inflation rates. Furthermore, the point forecasts made by the individual SPF forecasters with ID 463 (magenta line) and ID 483 (cyan line) are shown. Note that several of the forecasts are identical from early 2010 on.

2010, Section 3.2), and we refer to it as the SPF median with MSE technique.

Tables 4 and 5 compare the predictive performance of the postprocessing methods to that of the original SPF and probabilistic no-change forecasts, in terms of the MAE and the mean CRPS. In nowcasting at a prediction horizon of one quarter, none of the post-processing techniques outperforms the SPF forecast, neither in terms of the MAE nor the mean CRPS. At prediction horizons of two to five quarters, original and postprocessed SPF forecasts as well as the probabilistic no-change forecast all yield about equal MAE. Post-processing generally improves the mean CRPS, reaching a predictive performance that is comparable, but not superior, to that of the probabilistic no-change forecast.

Table 4: Mean absolute error (MAE) for forecasts of United States inflation from the third quarter of 1995 to the first quarter of 2010, in percentage points, along with the lower tail probability of the Diebold-Mariano test under the hypothesis of the predictive performance being equal to that of the SPF forecast. The probabilistic no-change forecast (PNC) is obtained using 20 quarters of data, and the length of the training period for all postprocessing methods is 40 quarters.

Forecast method	Prediction horizon in quarters				
	1	2	3	4	5
SPF	0.89	1.44	1.51	1.49	1.49
Probabilistic no-change	1.45 ₉₉	1.46 ₅₈	1.45 ₁₈	1.48 ₃₉	1.48 ₄₄
HR with SPF covariates	0.91 ₆₀	1.42 ₃₆	1.50 ₄₄	1.53 ₆₃	1.49 ₅₀
HR with SPF and PNC covariates	0.87 ₄₄	1.40 ₂₇	1.51 ₄₉	1.47 ₄₀	1.48 ₅₁
GM	1.05 ₉₉	1.43 ₄₇	1.45 ₁₈	1.47 ₃₆	1.47 ₃₇
GM with variance adjustment	0.89 ₉₁	1.47 ₆₉	1.50 ₃₄	1.48 ₃₉	1.45 ₁₆

Table 5: Mean continuous ranked probability score (CRPS) for forecasts of United States inflation from the third quarter of 1995 to the first quarter of 2010, in percentage points, along with the lower tail probability of the Diebold-Mariano test under the hypothesis of the predictive performance being equal to that of the SPF forecast. The probabilistic no-change forecast (PNC) is obtained using 5 years of data and the length of the training period for all postprocessing methods is 40 quarters.

Forecast method	Prediction horizon in quarters				
	1	2	3	4	5
SPF	0.69	1.16	1.25	1.26	1.27
SPF median with MSE	0.70 ₅₃	1.10 ₃₉	1.15 ₃₆	1.14 ₃₄	1.13 ₃₁
Probabilistic no-change	1.08 ₉₉	1.10 ₂₁	1.10 ₀₀	1.10 ₀₀	1.11 ₀₀
HR with SPF covariates	0.69 ₄₈	1.08 ₀₉	1.15 ₀₈	1.15 ₁₂	1.13 ₁₀
HR with SPF and PNC covariates	0.68 ₄₀	1.09 ₁₈	1.15 ₁₂	1.14 ₁₁	1.12 ₁₂
GM	0.80 ₉₆	1.09 ₁₇	1.10 ₀₁	1.11 ₀₁	1.10 ₀₀
GM with variance adjustment	0.68 ₃₆	1.10 ₂₀	1.11 ₀₀	1.12 ₀₀	1.10 ₀₀




Figure 6: MAE (solid lines) and mean CRPS (dashed lines) for forecasts of United States inflation from the first quarter of 2000 to the first quarter of 2010, as a function of the length of the rolling training period, at prediction horizons of one quarter (left) and five quarters (right). Results are shown for the SPF (blue), probabilistic no-change (red), HR with SPF covariates (yellow), HR with SPF and PNC covariates (orange), GM (dark green), and GM with variance adjustment (pale green) techniques.

4 Robustness

We have argued that the SPF median forecast outperforms simple no-change forecasts of United States inflation in current quarter nowcasts. However, at prediction horizons from two to five quarters ahead, probabilistic no-change forecasts have equal or higher skill than the SPF forecast, even after postprocessing. In making such claims, it is critically important to demonstrate the robustness of the results under changes in the details of the prediction experiment.

An initial robustness check was done in Section 2, where Figure 2 showed the performance of the probabilistic no-change forecast to depend little on the choice of the length of the rolling training period, thereby justifying that we fix it at 20 quarters. In this section we show that our key findings remain valid under changes in the length of the training period

for the postprocessing techniques, and we check whether they hold in smaller test periods, under distinct economic regimes.

Figure 6 plots the MAE and the mean CRPS for our postprocessing methods in their dependence on the length of the rolling training period, at prediction horizons of one and five quarters ahead. Except for the probabilistic no-change forecast, all methods show substantially higher skill at the shorter prediction horizon. Furthermore, for all methods the mean CRPS is much lower than the MAE. The predictive skill of the postprocessing methods does not depend much on the length of the training period, and our choice of a 40 quarter period seems reasonable.

Next we assess the effect of the test period and the corresponding economic regimes. In Tables 6 and 7, the aggregate results in Tables 1 and 2, which cover the third quarter of 1995 through the first quarter of 2010, have been stratified into three sub-periods of about equal length. With shorter test periods, the Diebold-Mariano test statistic is occasionally ill defined because of a negative variance estimate. Diebold and Mariano (1995) suggest that the variance estimate should then be treated as zero and the null hypothesis of equal forecast accuracy be rejected.

The first sub-period ranges from the third quarter of 1995 through the fourth quarter of 2000. This was an era of general economic boom in the United States and inflation rates were particularly stable, which facilitated forecasting and is mirrored in low MAEs. Hence, the traditional no-change forecast performed quite well during this period, particularly at prediction horizons of three and four quarters, where it had lower MAE than all other forecasts, including the SPF and the probabilistic no-change forecasts. Another look at Figure 1 supports this choice, because the behavior of inflation rates in the late 1990s appears generally compatible with a symmetric random walk model, and the traditional no-change forecast is the Bayes predictor under such an assumption (Granger, 1969; Gneiting, 2010). Thus, during this period of sustained economic growth the choice of the traditional

no-change forecast as reference standard was appropriate.

The second sub-period extends from the first quarter of 2001 through the fourth quarter of 2005. The expansion that began in the early 1990s came to an end at the beginning of this period with the early 2000s recession, and inflation rates started to behave much more erratically, not unlike a white noise model. The erratic behavior becomes even more pronounced during the third sub-period, which ranges from the first quarter of 2006 through the first quarter of 2010. This includes the late 2000s recession, a financial crisis much more severe than that of the previous period. Its effect culminated in the extreme deflation observed in the fourth quarter of 2008, at -9.2 percentage points, and we now see substantially lower levels of predictability than before.

Generally, from 2000 on the probabilistic no-change forecast succeeds the traditional no-change forecast as the appropriate reference forecast. In line with the aggregate results in Tables 1 and 2, at prediction horizons from two to five quarters the SPF forecast tends to benefit from postprocessing in terms of the mean CRPS, but not in terms of the MAE. We conclude that the spread adjustment generally is useful, but not the location adjustment. However, while the probabilistic forecast performance improves under postprocessing as measured by the mean CRPS, the postprocessed forecasts are unable to outperform the simplistic probabilistic no-change forecast.

The sub-period from 2006 on also is the longest consecutive period for which a complete set of forecasts is available for more than one individual SPF expert. In Table 1 we show MAEs for the five professionals who had complete records during this period. The predictive performance of the experts with IDs 463 and 483 is particularly impressive, and their point forecasts are illustrated in Figure 5.

Table 6: Mean absolute error (MAE) for forecasts of United States inflation from the third quarter of 1995 to the fourth quarter of 2000 (top panel), from the first quarter of 2001 to the fourth quarter of 2005 (middle panel), and from the first quarter of 2006 to the first quarter of 2010 (bottom panel), in percentage points, along with the lower tail probability of the Diebold-Mariano test under the hypothesis of the predictive performance being equal to that of the SPF forecast. The NA symbol indicates a missing value due to a negative variance estimate.

Forecast method	Prediction horizon in quarters				
	1	2	3	4	5
1995:Q3 - 2000:Q4					
SPF	0.52	0.87	0.93	0.95	0.95
Probabilistic no-change	0.87 ₉₉	0.90 ₇₂	0.93 ₄₉	0.96 _{NA}	0.98 _{NA}
Traditional no-change	0.76 ₉₉	0.92 ₆₆	0.78 ₁₄	0.88 ₃₃	1.15 ₉₄
HR with SPF covariates	0.54 ₆₅	0.84 ₃₇	0.94 ₅₂	1.00 ₅₇	0.99 ₅₅
HR with SPF and PNC covariates	0.56 ₇₅	0.88 ₅₆	0.96 ₆₀	1.03 ₆₂	1.01 ₅₈
GM	0.59 ₉₉	0.91 ₉₀	0.94 ₆₇	0.96 _{NA}	0.97 _{NA}
GM with variance adjustment	0.52 ₉₉	0.87 ₉₂	0.93 ₁₇	0.95 _{NA}	0.95 ₉₁
2001:Q1 - 2005:Q4					
SPF	1.00	1.26	1.33	1.26	1.28
Probabilistic no-change	1.18 ₈₇	1.18 ₁₂	1.17 ₁₀	1.16 ₂₄	1.13 ₀₁
Traditional no-change	1.65 ₉₉	1.52 ₉₀	1.74 ₉₇	1.86 ₉₉	1.58 ₉₉
HR with SPF covariates	1.11 ₉₈	1.20 ₀₃	1.25 ₂₂	1.22 ₃₆	1.18 ₀₀
HR with SPF and PNC covariates	1.13 ₉₉	1.21 ₁₁	1.25 ₂₅	1.22 ₃₇	1.17 ₀₀
GM	1.08 ₈₃	1.17 ₁₀	1.16 ₀₈	1.15 ₂₀	1.14 ₀₁
GM with variance adjustment	1.00 ₈₃	1.24 ₀₀	1.29 ₀₉	1.19 ₁₃	1.17 ₁₁
2006:Q1 - 2010:Q1					
SPF	1.22	2.39	2.48	2.47	2.45
Probabilistic no-change	2.52 ₉₉	2.51 ₆₅	2.46 ₄₅	2.52 ₇₀	2.55 ₆₉
Traditional no-change	3.36 ₉₉	4.18 ₉₆	3.88 ₉₀	3.84 ₉₂	3.72 ₉₄
HR with SPF covariates	1.14 ₃₈	2.43 ₆₀	2.53 ₆₆	2.59 ₇₈	2.52 ₇₇
HR with SPF and PNC covariates	1.00 ₂₀	2.33 ₃₈	2.51 ₅₇	2.35 ₀₀	2.51 ₆₂
GM	1.62 ₉₇	2.43 ₅₆	2.46 ₄₅	2.52 ₆₉	2.52 ₆₅
GM with variance adjustment	1.22 ₈₄	2.52 ₇₃	2.48 ₅₀	2.52 ₇₁	2.44 ₄₈
SPF forecaster ID 463	1.06 ₁₆	2.35 ₃₄	2.57 ₇₅	2.41 ₁₃	2.33 ₀₀
SPF forecaster ID 483	2.12 ₉₉	2.25 ₂₂	2.30 ₁₅	2.26 ₀₅	2.33 ₁₉
SPF forecaster ID 510	2.32 ₉₉	2.67 ₉₈	2.48 ₅₀	2.71 _{NA}	2.55 ₉₆
SPF forecaster ID 516	1.53 ₈₃	3.34 ₉₉	3.03 _{NA}	2.78 ₉₀	2.94 ₉₉
SPF forecaster ID 528	1.40 ₆₈	2.59 ₇₆	3.00 ₉₈	2.70 _{NA}	2.55 ₉₀

Table 7: Mean continuous ranked probability score (CRPS) for forecasts of United States inflation from the third quarter of 1995 to the fourth quarter of 2000 (top panel), from the first quarter of 2001 to the fourth quarter of 2005 (middle panel), and from the first quarter of 2006 to the first quarter of 2010 (bottom panel), in percentage points, along with the lower tail probability of the Diebold-Mariano test under the null hypothesis of equal predictive performance between the SPF forecast and the reference forecast. The NA symbol indicates a missing value due to a negative variance estimate.

Forecast method	Prediction horizon in quarters				
	1	2	3	4	5
1995:Q3 - 2000:Q4					
SPF	0.39	0.66	0.73	0.76	0.77
SPF median with MSE	0.39 ₃₈	0.59 ₀₄	0.65 ₁₂	0.71 ₂₄	0.72 ₃₀
Probabilistic no-change	0.59 ₉₉	0.62 ₂₀	0.63 ₀₁	0.66 ₀₀	0.68 ₀₁
Traditional no-change with MSE	0.55 ₉₈	0.63 ₃₆	0.57 ₁₀	0.70 ₃₁	0.86 ₈₁
HR with SPF covariates	0.40 ₆₂	0.60 ₂₁	0.66 ₂₆	0.73 ₄₄	0.74 ₄₄
HR with SPF and PNC covariates	0.42 ₇₆	0.63 ₃₄	0.68 ₃₀	0.75 ₄₈	0.76 ₄₈
GM	0.43 ₉₈	0.63 ₂₀	0.66 ₀₁	0.67 ₀₀	0.69 ₀₀
GM with variance adjustment	0.38 ₂₂	0.58 ₀₀	0.61 ₀₀	0.64 ₀₀	0.66 ₀₂
2001:Q1 - 2005:Q4					
SPF	0.82	1.03	1.07	1.07	1.07
SPF median with MSE	0.82 ₅₉	0.94 ₀₁	0.97 ₀₀	0.93 ₀₀	0.92 ₀₀
Probabilistic no-change	0.87 ₇₃	0.88 ₀₁	0.89 ₀₀	0.89 ₀₀	0.87 ₀₀
Traditional no-change with MSE	1.17 ₉₈	1.20 ₈₆	1.26 ₈₃	1.41 ₉₆	1.15 ₇₂
HR with SPF covariates	0.85 ₇₈	0.93 ₀₂	0.97 ₀₇	0.94 ₀₆	0.88 ₀₀
HR with SPF and PNC covariates	0.88 ₉₀	0.94 ₀₃	0.98 ₀₉	0.95 ₀₆	0.87 ₀₀
GM	0.81 ₄₄	0.88 ₀₁	0.86 ₀₀	0.86 ₀₂	0.83 ₀₀
GM with variance adjustment	0.80 ₂₁	0.92 ₀₀	0.94 ₀₀	0.91 ₀₀	0.85 ₀₀
2006:Q1 - 2010:Q1					
SPF	0.92	1.98	2.13	2.15	2.17
SPF median with MSE	0.95 ₇₁	1.96 ₄₆	2.02 ₃₂	1.96 _{NA}	1.90 _{NA}
Probabilistic no-change	1.95 ₉₉	1.98 ₅₀	1.94 ₀₉	1.95 ₀₈	1.95 _{NA}
Traditional no-change with MSE	3.31 ₉₉	3.54 ₉₈	3.03 ₈₄	2.87 ₈₅	2.64 ₇₉
HR with SPF covariates	0.86 ₃₉	1.87 ₂₉	1.99 ₂₂	1.95 ₁₁	1.95 _{NA}
HR with SPF and PNC covariates	0.78 ₂₁	1.86 ₃₁	1.97 ₂₄	1.91 ₁₃	1.91 _{NA}
GM	1.26 ₉₇	1.93 ₄₂	1.96 ₁₅	1.96 ₁₅	1.95 ₁₂
GM with variance adjustment	0.93 ₅₈	1.99 ₅₂	1.96 ₁₃	1.99 ₁₂	1.95 ₀₇

5 Conclusions

Predicting inflation is important, and there are various ways of doing it, including forecasts from the Survey of Professional Forecasters (SPF) and simple no-change forecasts. In the extant literature, the traditional no-change forecast has served as a benchmark, to which more sophisticated techniques are to be compared. While this was appropriate during the late 1990s economic boom, in today’s turbulent markets the equally simplistic probabilistic no-change forecast performs much better. To avoid spurious claims of predictability, we make a plea for the use of the probabilistic no-change forecast as a default reference standard in inflation forecasting.

In current quarter nowcasting, corresponding to a prediction horizon of one quarter, economic expertise is persuasive, and the SPF professionals outperform all types of no-change forecasts. Businesses, organizations and the government are well advised to avail themselves of the SPF experts’ short-term predictions.

At prediction horizons beyond the current quarter, the probabilistic predictive performance of the SPF forecast, as measured by the mean CRPS, improves under statistical postprocessing. Supplementing the SPF median with MSE method, which has been used by the Federal Reserve Bank of Philadelphia, we have introduced heteroscedastic regression (HR) and Gaussian mixture (GM) techniques for doing this. However, at prediction horizons of two or more quarters, even postprocessed SPF forecasts fail to outperform the simplistic probabilistic no-change forecast.

While novel and potentially surprising in the specific context of the SPF and inflation rates, the result conforms with a general theme in the forecasting literature, in that simple prediction methods tend to perform well, with overfitting being heavily penalized and subjective human expertise often being overrated. For example, Nelson (1972) showed that simple statistical methods can outperform complex economic simulation models, and a recent 20-year study argues persuasively that professional forecasters, who appear as ex-

perts in the media and advise governments, organizations and businesses, might not be better prognosticians than John Q. Public (Tetlock, 2005). In this light, we might be well advised to adapt to low levels of socio-economic predictability, where uncertainty reigns (Makridakis and Taleb, 2009).

Appendix

Here we provide a closed form expression for the continuous ranked probability score (1) when the predictive distribution is the three-parameter two-piece normal distribution with CDF given by (2). A tedious but straightforward calculation shows that

$$\text{CRPS}(F_{\text{TPN}}, y) = \begin{cases} \frac{4\sigma_1^2}{\sigma_1 + \sigma_2} \left[\frac{y - \mu}{\sigma_1} \Phi\left(\frac{y - \mu}{\sigma_1}\right) + \varphi\left(\frac{y - \mu}{\sigma_1}\right) \right] \\ \quad - (y - \mu) \\ \quad + \frac{2}{\sqrt{\pi}} \frac{\sqrt{2}\sigma_2(\sigma_2^2 - \sigma_1^2) - (\sigma_1^3 + \sigma_2^3)}{(\sigma_1 + \sigma_2)^2} & \text{if } y \leq \mu, \\ \frac{4\sigma_2^2}{\sigma_1 + \sigma_2} \left[\frac{y - \mu}{\sigma_2} \Phi\left(\frac{y - \mu}{\sigma_2}\right) + \varphi\left(\frac{y - \mu}{\sigma_2}\right) \right] \\ \quad + \frac{(\sigma_1 - \sigma_2)^2 - 4\sigma_2^2}{(\sigma_1 + \sigma_2)^2} (y - \mu) \\ \quad + \frac{2}{\sqrt{\pi}} \frac{\sqrt{2}\sigma_1(\sigma_1^2 - \sigma_2^2) - (\sigma_1^3 + \sigma_2^3)}{(\sigma_1 + \sigma_2)^2} & \text{if } y \geq \mu. \end{cases} \quad (8)$$

Grimmett et al. (2006) give a similar formula when the predictive distribution is a mixture of Gaussian components, which we use to compute the CRPS for the GM forecasts.

Acknowledgments

We are grateful to Tom Stark for generously sharing his knowledge about the SPF data, and to Marc Genton and Alex Lenkoski for discussions. Funding was provided by NSF grant

DMS-0706745 to the University of Washington and by the Alfried Krupp von Bohlen und Halbach-Foundation.

References

Ang, A., Bekaert, G., and Wei, M. “Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?” *Journal of Monetary Economics*, 54:1163–1212 (2007).

Croushore, D. “Introducing: The Survey of Professional Forecasters.” *Federal Reserve Bank of Philadelphia Business Review*, 3–13 (November/December 1993).

Croushore, D. and Stark, T. “A Real-Time Data Set for Macroeconomists.” *Journal of Econometrics*, 105:111–130 (2001).

Diebold, F. X., Gunther, T. A., and Tay, A. S. “Evaluating Density Forecasts with Applications to Financial Risk Management.” *International Economic Review*, 39:863–883 (1998).

Diebold, F. X. and Mariano, R. S. “Comparing Predictive Accuracy.” *Journal of Business and Economic Statistics*, 13:253–263 (1995).

Elder, R., Kapetanios, G., Taylor, T., and Yates, T. “Assessing the MPC’s Fan Charts.” *Bank of England Quarterly Bulletin*, 326–348 (Autumn 2005).

Elliott, G. and Timmermann, A. “Economic Forecasting.” *Journal of Economic Literature*, 46:1–53 (2008).

Engelberg, J., Manski, C. F., and Williams, J. “Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters.” *Journal of Business and Economic Statistics*, 27:30–41 (2009).

- Gneiting, T. “Editorial: Probabilistic Forecasting.” *Journal of the Royal Statistical Society Series A*, 171:319–321 (2008).
- . “Making and Evaluating Point Forecasts.” Preprint (2010). Available online at <http://arxiv.org/abs/0912.0902>.
- Gneiting, T. and Raftery, A. E. “Weather Forecasting with Ensemble Methods.” *Science*, 310:248–249 (2005).
- . “Strictly Proper Scoring Rules, Prediction, and Estimation.” *Journal of the American Statistical Association*, 102:359–378 (2007).
- Gneiting, T., Raftery, A. E., Westveld, A. H., and Goldman, T. “Calibrated Probabilistic Forecasting Using Ensemble Model Output Statistics and Minimum CRPS Estimation.” *Monthly Weather Review*, 133:1098–1118 (2005).
- Granger, C. W. J. “Prediction with a Generalized Cost of Error Function.” *Operational Research Quarterly*, 20:199–207 (1969).
- Grimmett, E. P., Gneiting, T., Berrocal, V. J., and Johnson, N. A. “The Continuous Ranked Probability Score for Circular Variables and its Application to Mesoscale Forecast Ensemble Verification.” *Quarterly Journal of the Royal Meteorological Society*, 132:2925–2942 (2006).
- Hoeting, J. A., Madigan, D. M., Raftery, A. E., and Volinsky, C. T. “Bayesian Model Averaging: A Tutorial.” *Statistical Science*, 14:382–401 (1999).
- John, S. “The Three-Parameter Two-Piece Normal Family of Distributions and its Fitting.” *Communications in Statistics — Theory and Methods*, 11:879–885 (1982).
- Leslie, D. S., Kohn, R., and Nott, D. J. “A General Approach to Heteroscedastic Linear Regression.” *Statistics and Computing*, 17:131–146 (2007).

- Makridakis, S. and Taleb, N. “Living in a World of Low Levels of Predictability.” *International Journal of Forecasting*, 25:840–844 (2009).
- Matheson, J. E. and Winkler, R. L. “Scoring Rules for Continuous Probability Distributions.” *Management Science*, 22:1087–1096 (1976).
- Nelson, C. E. “The Prediction Performance of the FRB-MIT-PENN Model of the U.S. Economy.” *American Economic Review*, 62:902–917 (1972).
- Palmer, T. N. “The Economic Value of Ensemble Forecasts as a Tool for Risk Assessment: From Days to Decades.” *Quarterly Journal of the Royal Meteorological Society*, 128:747–774 (2002).
- R Development Core Team. *R: a Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical Computing (2008).
- Raftery, A. E., Gneiting, T., Balabdaoui, F., and Polakowski, M. “Using Bayesian Model Averaging to Calibrate Forecast Ensembles.” *Monthly Weather Review*, 133:1155–1174 (2005).
- Sloughter, J. M., Gneiting, T., and Raftery, A. E. “Probabilistic Wind Forecasting Using Ensembles and Bayesian Model Averaging.” *Journal of the American Statistical Association*, 105:25–35 (2010).
- Stark, T. “Error statistics for the Survey of Professional Forecasters for QA/QA CPI Inflation [AR, PPs].” (2010). Available online at http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/data-files/CPI/SPF_Error_Statistics_CPI_1_AIC.pdf.
- Tetlock, P. E. *Expert Political Judgement: How Good Is It? How Can We Know?*. Princeton University Press (2005).

- Thorarinsdottir, T. L. and Gneiting, T. “Probabilistic Forecasts of Wind Speed: Ensemble Model Output Statistics by Using Heteroscedastic Censored Regression.” *Journal of the Royal Statistical Society Series A*, 173:371–388 (2010).
- Timmermann, A. “Density Forecasting in Economics and Finance.” *Journal of Forecasting*, 19:231–234 (2000).
- Wallis, K. F. “Asymmetric Density Forecasts of Inflation and the Bank of England’s Fan Chart.” *National Economic Review*, 167:106–112 (1999).
- Zarnowitz, V. “The New ASA-NBER Survey of Forecasts by Economic Statisticians.” *American Statistician*, 23:12–16 (1969).